Pergamon

PII: S0042-6989(96)00064-8

Vision Res., Vol. 36, No. 19, pp. 3061–3075, 1996
Published by Elsevier Science Ltd.
Printed in Great Britain
0042-6989/96 $15.00 + 0.00

The Barberplaid Illusion: Plaid Motion is Biased by Elongated Apertures

BRENT R. BEUTTER,* JEFFREY B. MULLIGAN,* LELAND S. STONE*†

Received 28 August 1995; in final form 1 February 1996

The perceived direction of motion of plaids windowed by elongated spatial Gaussians is biased toward the window's long axis. The bias increases as the relative angle between the plaid motion and the long axis of the window increases, peaks at a relative angle of ~45 deg, and then decreases. The bias increases as the window is made narrower (at fixed height) and decreases as the component spatial frequency increases (at fixed aperture size). We examine several models of human motion processing (cross-correlation, motion-energy, intersection-of-constraints, and vector-sum), and show that none of these standard models can predict our data. We conclude that spatial integration of motion signals plays a crucial role in plaid motion perception and that current models must be explicitly expanded to include such spatial interactions. Published by Elsevier Science Ltd.

Plaids Motion models Direction discrimination Aperture Intersection-of-constraints rule

INTRODUCTION

Human perception of motion depends not only on the physical motion of objects, but also on the conditions under which the motion is viewed. A simple yet dramatic example of this is the barberpole illusion, in which the perceived direction of motion of an obliquely oriented drifting grating is vertical, when viewed through a vertically oriented rectangular aperture; but the same motion appears horizontal, when the rectangular aperture is horizontal. The influence of the aperture on the perceived motion of one-dimensional patterns has been extensively examined (e.g. Wallach, 1935; Mulligan, 1991; Power & Mouldon, 1992; Kooi, 1993; Mulligan & Beutter, 1994). The question that we examine in this paper is whether the perceived direction of motion of two-dimensional patterns such as plaids is also affected by the type of viewing window. We address this question by measuring the perceived direction of moving plaids windowed by elongated spatial Gaussians.

The direction of motion of a one-dimensional pattern viewed through a restricted aperture is inherently ambiguous, as shown in Fig. 1(a). The motion is consistent with any velocity vector that falls on the constraint line, because only the velocity component in the direction perpendicular to the orientation of the pattern is uniquely determined. Thus, one might expect that the type of window in which one-dimensional patterns are moved would determine how this ambiguity is resolved, and thus which of the many possible directions of motion is actually perceived. On the other hand, the direction of motion of two-dimensional patterns, such as plaids, is unambiguous (Adelson & Movshon, 1982), because the multiple constraints provided by the components allow only a single solution [see Fig. 1(b)]. The intersection-of-constraints rule specifies how this unique resultant pattern velocity could be computed from the component velocities and orientations (Fennema & Thompson, 1979; Adelson & Movshon, 1982). Is the human visual system able to extract this unique correct velocity independent of the aperture, or does the window shape affect the processing of the motion signals?

In primates, evidence from anatomy, physiology, and psychophysics suggests that motion processing of two-dimensional velocity appears to occur in a two-stage process. First, directionally selective mechanisms detect the motion of local one-dimensional features in the image (Hubel & Wiesel, 1968; Watson *et al.*, 1980; De Valois *et al.*, 1982b) and then a second stage integrates these one-dimensional signals over a larger area of the visual field to extract the two-dimensional pattern velocity (Adelson & Movshon, 1982; Maunsell & Van Essen, 1983a,b; Albright, 1984; Mikami *et al.*, 1986a,b; Movshon *et al.*, 1986, 1988; Newsome *et al.*, 1986; Ungerleider & Desimone, 1986; Rodman & Albright, 1989; Welch, 1989; Stone, 1990). In primate cerebral cortex, the first

*MS 262-2, Human and Systems Technologies Branch, Flight Management and Human Factors Division, NASA Ames Research Center, Moffett Field CA 94035-1000, U.S.A.
†To whom reprint requests should be addressed.

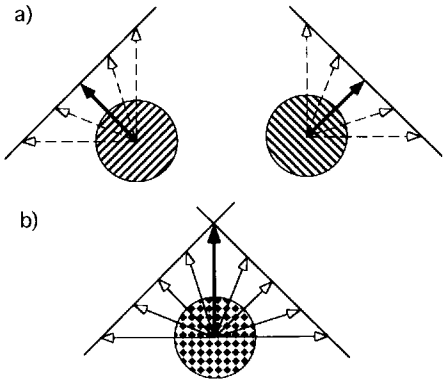

FIGURE 1. (a) Two one-dimensional gratings oriented ±45 deg from the vertical are shown in circular apertures. The velocity component in the direction perpendicular to the orientation of the grating is fixed by the constraint line (solid diagonal lines), while the component in the parallel directions is ambiguous. Velocities consistent with the constraint line are shown by the arrows. The velocity that is usually perceived is shown by the solid filled arrow. (b) When these two one-dimensional gratings are added together, the result is a two-dimensional plaid pattern, which has a unique velocity. The constraint lines for each grating intersect at a single point, which determines the unique IOC velocity of the plaid (solid filled arrow). The unfilled arrows which represent possible velocities of the one-dimensional gratings in (a), satisfy the constraint line of only one of the component gratings, but are inconsistent with the additional constraint imposed by the other grating, and therefore do not correspond to possible plaid velocities.

stage of motion processing occurs in primary visual cortex (V1). Motion-sensitive neurons in V1 have small receptive fields tuned to stimuli of a specific size, orientation, and direction of motion (Hubel & Wiesel, 1968; De Valois *et al.*, 1982b). Each of these cells responds maximally to the component of motion in the direction perpendicular to its preferred orientation and shows little or no response to motion parallel to its preferred orientation. Thus, when an object moves, V1 neurons respond to the local motion of one-dimensional features in the image, and therefore cannot individually signal the velocity of the full two-dimensional pattern. However, the actual two-dimensional pattern velocity can be recovered by combining the one-dimensional signals from multiple V1 neurons. Directionally selective V1 neurons project to the middle temporal (MT) area (Maunsell & Van Essen, 1983a; Movshon & Newsome, 1984), where most neurons appear to respond preferentially to motion (Zeki, 1974; Maunsell & Van Essen, 1983b; Albright, 1984; Albright *et al.*, 1984). Neurons in MT have larger receptive fields (Maunsell & Van Essen, 1983b) and there is some evidence that MT neurons integrate the one-dimensional edge motion signals to compute the two-dimensional pattern velocity (Movshon *et al.*, 1986; Rodman & Albright, 1989; Britten *et al.*, 1993). The above notwithstanding, there is also some psychophysical evidence for noncomponent driven one-stage velocity estimation [e.g. Derrington & Badcock (1992)]. Additionally, there are clearly other motion processing pathways leading to MT, directly from cortical areas V2 and V3 and from the Superior Colliculus via the pulvinar, which may also play an important role in velocity estimation (Maunsell & Van Essen, 1983a; Ungerleider *et al.*, 1984; De Yoe & Van Essen, 1985; Ungerleider & Desimone, 1986; Rodman *et al.*, 1989, 1990).

Modelers have also explored the ways primate visual cortex might estimate velocity. These models fall into several classes. Bulthoff *et al.* (1989) have proposed that velocity is estimated by finding the maximum of an image cross-correlation function. Motion-energy type models (Watson & Ahumada, 1983; van Santen & Sperling, 1984, 1985; Adelson & Bergen, 1985; Watson & Ahumada, 1985; Heeger, 1987) determine perceived velocity using more biologically plausible processes. First, the image is decomposed into its spatio-temporal components (much like what is done in V1). Then, velocity is estimated by finding the single pattern velocity most consistent with the entire motion-energy spectrum. Intersection-of-constraints (IOC) models (Fennema & Thompson, 1979; Adelson & Movshon, 1982) are also explicitly two-staged. First, the motions of the component gratings are estimated separately, and then the IOC rule is used to compute the pattern velocity that is consistent with all of the component constraints (Fig. 1). Chubb and Sperling (1988, 1989) first proposed the existence of two motion pathways: a Fourier pathway, operating directly on the stimulus, and a non-Fourier pathway that contains a nonlinear preprocessing stage, that performs a rectification or squaring prior to the motion processing. Wilson *et al.* (1992) formulated a motion model incorporating both these pathways. Their vector-sum model measures both Fourier and non-Fourier motion in separate pathways, and then combines these signals using a vector-sum rule to compute the direction of motion of the pattern.

In this paper, we first describe the effect of aperture shape on the perceived direction of moving two-dimensional patterns (plaids). We then use these results to determine if any of the above models can predict human performance for this type of plaid stimulus. Preliminary results have been presented elsewhere (Beutter *et al.*, 1994a,b).

METHODS

Observers

Four observers between the ages of 33 and 40yr participated in the experiments. One observer, PS, was naive but also strabismic.

Stimuli and apparatus

The stimulus, $I(\vec{x}, t)$, in these experiments was a drifting plaid, $P(\vec{x}, t)$, windowed spatially by a stationary

elongated Gaussian, $W(\vec{x})$, and temporally by a trapezoidal function, $H(t)$, as described by Eqs (1)–(4).

$$I(\vec{x},t) = I_0[1 + cP(\vec{x},t)W(\vec{x})H(t)] \quad (1)$$

where

$$P(\vec{x},t) = \sin\left[2\pi\left(\vec{f}_s \cdot \vec{x} + f_t t\right)\right] + \sin\left[2\pi\left(\vec{f}_s^{\perp} \cdot \vec{x} + f_t t\right)\right] \quad (2)$$

$$W(\vec{x}) = \exp\left(-\frac{(\vec{x}\cdot\vec{e})^2}{2\sigma_H^2} - \frac{(\vec{x}\cdot\vec{e}^{\perp})^2}{2\sigma_W^2}\right) \quad (3)$$

$$H(t) = \begin{cases} t/T_1, & 0 < t < T_1 \\ 1 & T_1 < t < T_2 \\ 1 - (t - T_2)/T_1 & T_2 < t < T_2 + T_1 \end{cases} \quad (4)$$

The plaid, $P(\vec{x},t)$, was the sum of two orthogonal ($\vec{f}_s \cdot \vec{f}_s^{\perp} = 0$) sine-wave "component" gratings moving with equal speeds. Both sine-wave gratings were of equal spatial frequencies, ($|\vec{f}_s| = |\vec{f}_s^{\perp}| = 0.3$, 0.6, or 1.2 c/deg), equal temporal frequencies ($f_t = 4$ Hz), and equal peak contrast ($c = 0.125$). The mean luminance, I_0, was fixed at 42 cd/m^2. Because the plaid was symmetric and its component speeds were equal, the plaid direction of motion was always midway between the orientations of the component gratings. We varied its direction of motion by rotating both component gratings equally. The spatial window, $W(\vec{x})$, was an elongated Gaussian with unequal standard deviations, σ_H (height) and σ_W (width), in its two principal directions, $\vec{e}$ and $\vec{e}^{\perp}$, respectively. In all experiments, σ_H was fixed at 2.5 deg, and the aspect ratio, σ_H/σ_W, was varied by setting σ_W to be 0.625, 1.25, or 1.77 deg. We defined the absolute window orientation to be the orientation of the unit vector $\vec{e}$. The time course of the stimulus was controlled by $H(t)$. It linearly ramped on for 167 msec (T_1), remained constant for 500 msec ($T_2 - T_1$), and then linearly ramped off for 167 msec. An example of a single of frame of the stimulus is shown in Fig. 2.

The stimuli were displayed on a 19" Barco color monitor (model CDCT 6351B) using an AT Vista video display system hosted by an IBM 486. The monitor was run in the 60-Hz interlaced mode. To minimize interlace artifacts, alternate horizontal lines were set equal to one another by computing a 320×243 pixel image and then zooming it by a factor of two in both the horizontal and vertical directions so that it filled the 640×486 display region. For the spatial and temporal frequencies of our stimuli, no significant aliasing occurred. The display pixel sizes were 0.47 mm horizontally and 0.54 mm vertically. At the 57-cm viewing distance, the full display subtended 30 deg × 26 deg. The luminance output of the monitor was calibrated to correct for its gamma nonlinearity using a lookup table.

The plaid motion was produced by using a dithering animation method which is described in detail in Mulligan and Stone (1989). Briefly, to generate a single drifting sinusoidal grating, it uses two sub-component gratings differing in spatial phase by 90 deg to produce a sum grating whose phase is varied by modifying the sub-component gratings' relative amplitudes through changes in the lookup table. The dithering procedure produces low-contrast spatial artifacts, which are near or below threshold (Mulligan & Stone, 1989). The plaid stimulus was constructed by creating four sub-component sinusoidal gratings (a sine and cosine sub-component pair for each plaid component) which were then multiplied by the elongated spatial Gaussian window function to produce four grating images. Each of these four images was then dithered to three levels using a modified error-diffusion algorithm (Mulligan, 1986). These images were then combined to produce a final image containing 81, (3^4), gray levels. For each frame, the appropriate lookup table values of these gray levels were precomputed and stored. Plaid motion was produced by sequentially loading the lookup tables on a frame-by-frame basis. For each combination of plaid and window angle, a separate image was created.

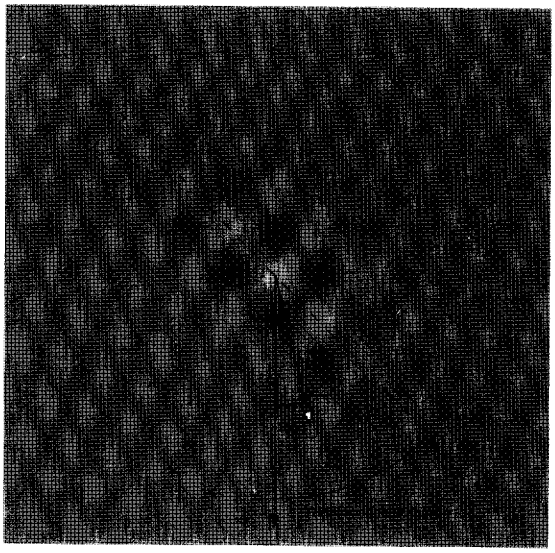

FIGURE 2. This figure shows a single frame of a two-dimensional plaid windowed by an elongated spatial Gaussian. The window has an aspect ratio of 2.0 and is oriented 40 deg to the right of straight down. The direction of the plaid motion is straight down as indicated by the filled arrow, while the perceived direction of motion is biased toward the window orientation, as indicated by the open arrow.

Procedure and data analysis

We used the method of adjustment to measure the perceived direction of plaid motion. After the stimulus had been presented, a pointer that subtended 15 deg appeared in the center of the screen. Two keys allowed the observer to rotate the pointer about the center of the display in either a clockwise or a counterclockwise direction. The observer was asked to adjust the orientation of the pointer so that it was aligned with the perceived direction of plaid motion. Although there was no limit to the time allowed to make the setting, observers

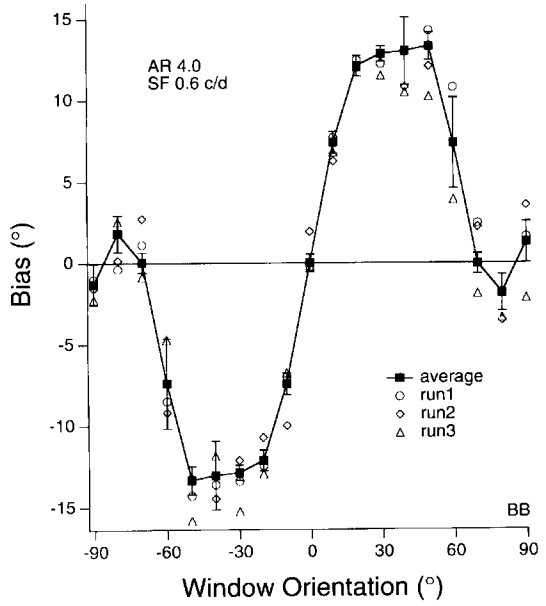

FIGURE 3. The biases in perceived direction are plotted as a function of the window orientation for observer BB. The stimulus had an aspect ratio of 4.0 and a component spatial frequency of 0.6 c/deg. The open symbols show the biases (8 settings) found in three separate runs for each window orientation, while the filled squares show the average biases. The average biases (48 settings) were computed by combining the data for positive and negative window orientations for each run, and then averaging over runs. Thus, the magnitudes of the biases for the positive and negative window orientations are identical. Both are presented for clarity in this and subsequent figures. The error bars represent the standard deviations over runs. In this and subsequent figures, the horizontal line represents zero bias.

generally did so in 1–3 sec. Before data collection began, observers were shown sample stimuli and practiced adjusting the pointer. No feedback was given.

We defined the window aspect ratio to be the ratio of the window's height to its width, σ_H/σ_W. Window orientation was defined relative to the plaid as the absolute window orientation minus the absolute direction of motion. Similarly, the bias was defined as the perceived direction of motion relative to the absolute plaid direction of motion. In a preliminary experiment, we verified that a circularly symmetric window (aspect ratio equal to 1.0) produced negligible biases.

For each run, the spatial frequency of the components and the aspect ratio of the Gaussian window were fixed. Runs consisted of 152 settings: two repetitions of all combinations of four plaid directions of motion (± 20 and ± 40 deg), and of 19 window orientations (0, ± 10, ± 20, ± 30, ± 40, ± 50, ± 60, ± 70, ± 80, ± 90 deg). Each observer performed a minimum of three runs in each experiment. Because the results for each absolute plaid direction were similar, the bias for each window orientation was computed by first averaging the results from the four plaid directions. Figure 3 shows the biases for observer BB, as a function of the window orientation for three separate runs of Experiment 2, with a window aspect ratio of 4.0 and a component spatial frequency of 0.6 c/deg. The results were symmetric about zero: positive window orientations produced biases approximately equal to, but opposite in sign to, those produced by negative window orientations of the same magnitude. Because of this symmetry, the results for the negative and positive window orientations for each run were combined to compute the mean bias for each window orientation and these were then averaged over runs to compute the average bias (filled squares). In subsequent figures, only the average biases are shown. The error bars are the standard deviations of the mean biases over the individual runs.

RESULTS

In Experiment 1, we measured the biases in the perceived direction of plaid motion produced by an elongated window, with an aspect ratio of 2, for a plaid composed of 0.6 c/deg sinusoidal gratings. Plots of the biases as a function of the relative window orientation for four observers are shown in Fig. 4. The results for all the observers were similar: they showed biases toward the long axis of the window. For window orientations <~45 deg, the biases increased as the window orientation increased. For window orientations >~45 deg, the biases decreased, and even reversed for one observer (PS). When the plaid moved in the direction of the short axis of the window (± 90 deg), the biases were negligible. The peak biases of the four observers ranged from 7.4 to 11.1 deg, and occurred at a window orientation of ~45 deg.

In Experiment 2, we investigated the effect of varying the window aspect ratio on the perceived direction of plaid motion. The biases produced by aspect ratios of 1.4, 2.0, and 4.0 for a plaid composed of 0.6 c/deg sinusoidal gratings are shown in Fig. 5. The results for the two observers tested were similar: the biases increased as the window aspect ratio increased. The pattern of results found for an aspect ratio of 2.0 was also found for the 1.4 and 4.0 aspect ratios: biases increased as the window orientation increased, reached a peak at an angle of about 45 deg, and then decreased. A summary of these results is shown as the filled circles in Fig. 8(a), in which the magnitude of the peak bias is plotted as a function of the aspect ratio. The largest aspect ratio, 4.0, produced the largest biases, while the smallest aspect ratio produced the smallest biases. These data show that long narrow windows (high aspect ratios) produce large biases, while more circular windows (low aspect ratios) produce smaller biases.

In Experiment 3, we examined the effects of varying the spatial frequency of the plaid components, while keeping the window constant. Because the window parameters were fixed, varying the spatial frequency also changed the number of visible cycles of the component gratings. For simplicity, we have chosen to discuss the data directly in terms of the component spatial frequency,

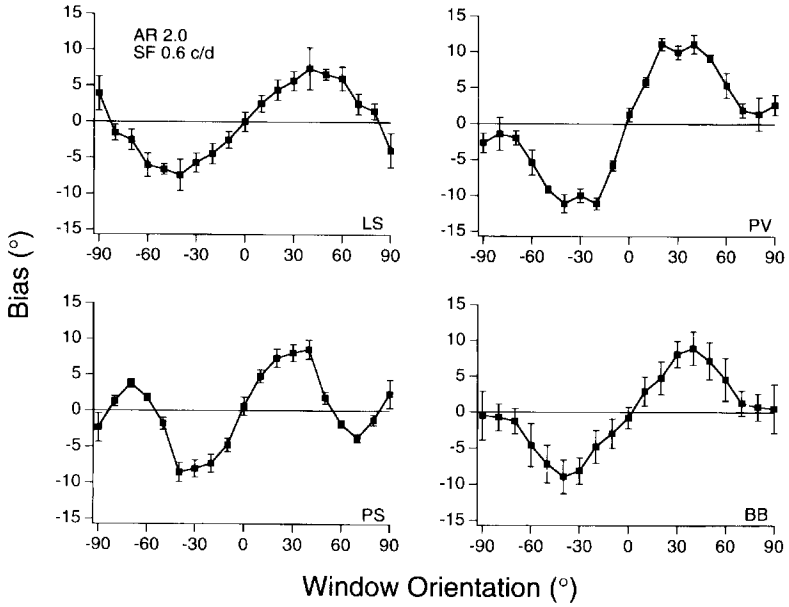

FIGURE 4. The average bias in perceived direction is plotted as a function of window orientation for a stimulus with an aspect ratio of 2.0 and component spatial frequency of 0.6 c/deg. Results are shown for four observers, including naive observer, PS. The error bars represent the standard deviation over runs.

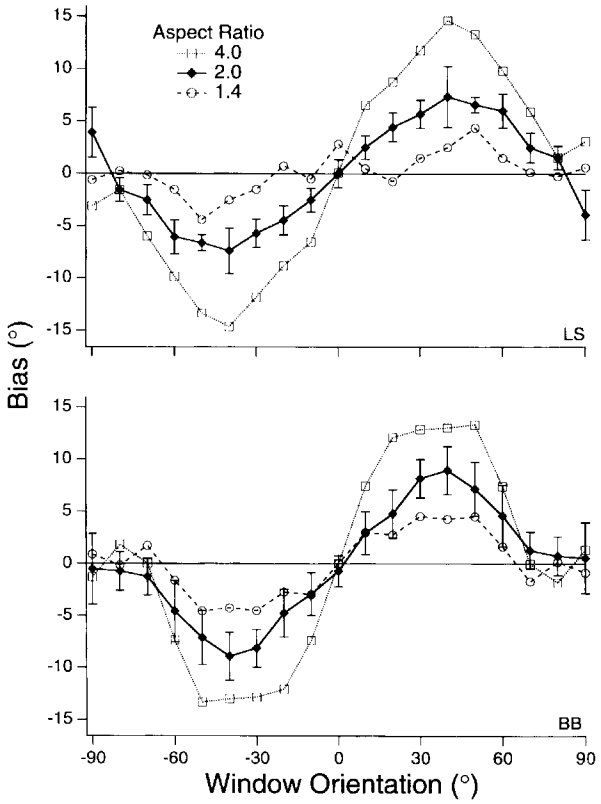

FIGURE 5. The average bias in perceived direction as a function of window orientation is plotted for stimuli with component spatial frequency fixed at 0.6 c/deg and three aspect ratios: 1.4 (○), 2.0 (◆), and 4.0 (□). Results are shown for two observers. For clarity, error bars are shown only for the aspect ratio equal to 2.0 data and represent the standard deviation over runs.

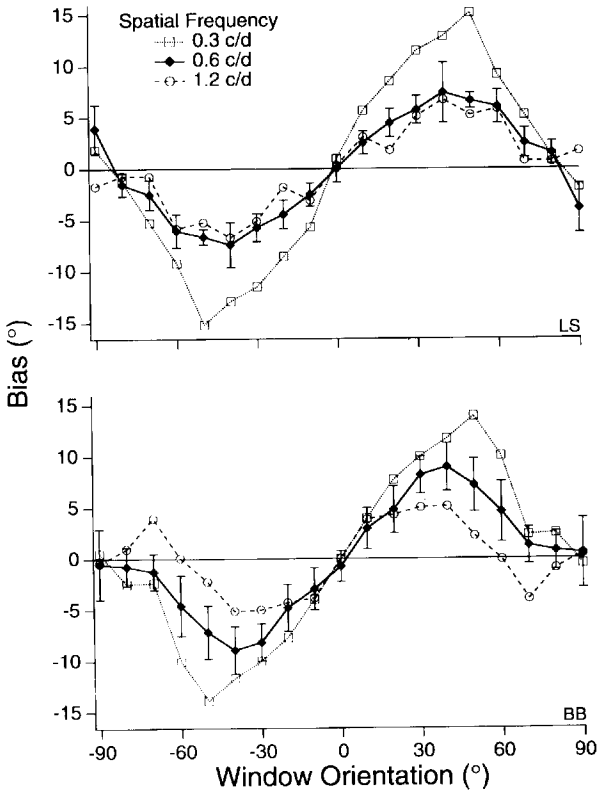

FIGURE 6. The average biases in perceived direction are plotted as a function of window orientation for stimuli with an aspect ratio fixed at 2.0 and three component spatial frequencies: 0.3 (□), 0.6 (◆), and 1.2 c/deg (○). Results are shown for two observers. For clarity, error bars are shown only for the 0.6 c/deg data and represent the standard deviation over runs.

rather than in terms of bandwidth or number of visible cycles. The biases for plaids composed of sinusoidal gratings with spatial frequencies of 0.3, 0.6, and 1.2 c/deg, and an aspect ratio of 2.0 are shown in Fig. 6. The results for the two observers were similar: the biases decreased as the component spatial frequency increased. The pattern of results found for each of the spatial frequencies was similar: biases increased as the window orientation increased, reached a peak at an angle of ~45 deg, and then decreased. A summary of these results is shown as the filled circles in Fig. 8(b), in which the magnitude of the peak bias is plotted as a function of the component spatial frequency. The lowest spatial frequency produced the largest biases, while the highest spatial frequency produced the smallest biases.

MODELING

Unwindowed moving plaids have an unambiguous direction of motion which can be found by using the IOC rule (Fig. 1). Our results however clearly show that plaids in elongated windows are not always seen to move in this direction. Rather, observers consistently report a bias in the perceived direction of plaid motion toward the direction of the long axis of the window. To understand how this bias may arise, we examined the predictions of several models of human motion processing: a cross-correlation model, modified IOC models, a motion-energy model, and Wilson, Ferrera and Yo's vector-sum model (1992).

Cross-correlation models

Cross-correlation models [e.g. Leese *et al.* (1970); Bulthoff *et al.* (1989)] determine the direction of motion by computing the translation that produces the maximal overlap between the image at two different times. We calculated the global cross correlation for plaids moving within elongated windows by using the following equation:

$$CC(\Delta x, \Delta y, t, \Delta t) = \int_{-\infty}^{\infty} dx \int_{-\infty}^{\infty} dy \cdot I(x, y, t) \cdot I(x + \Delta x, y + \Delta y, t + \Delta t) \quad (5)$$

We determined the predicted direction of motion by computing the velocity ($\Delta x/\Delta t$, $\Delta y/\Delta t$) that maximized the cross-correlation.* The predicted biases for an aspect ratio of 2.0 and component spatial frequency of 0.6 c/deg

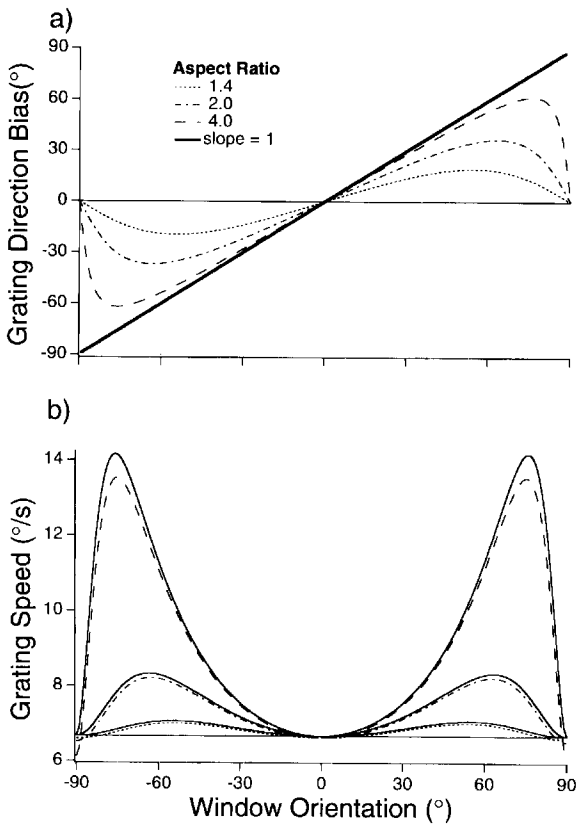

FIGURE 9. The biased component grating directions and speeds predicted by the cross-correlation calculation are shown for three aspect ratios: 1.4 (dotted lines), 2.0 (dot-dashed lines), and 4.0 (dashed lines). (a) The predicted biases in the grating direction of motion are plotted as a function of the window orientation. The thick diagonal line with slope equal to one represents the prediction that the perceived motion is in the direction of the window's orientation. (b) The predicted biases in the grating speed are plotted as a function of the window orientation. The solid line associated with the prediction for each aspect ratio, is the speed defined by the grating's constraint line and its direction bias. It is computed as $V_0/\cos(\theta_B)$ where θ_B is the bias in the perceived grating direction and V_0 is the unbiased grating speed (temporal frequency/spatial frequency).

dimensional pattern from the ambiguous motion of its one-dimensional component gratings (Fig. 1). Models implementing the IOC rule have two stages: a first stage measures the grating motion; and a second stage uses the constraints provided by the gratings to compute the resultant velocity.

In principle, to determine the constraint for each component, IOC models require *three* pieces of information: component orientation, speed, and direction of motion. Each constraint line's orientation is parallel to the component's orientation and its location is determined by the component's speed and direction. To generate a full set of exact IOC predictions from the biases in the perception of the component gratings, therefore would require explicit psychophysical measurements of each of these three quantities, for all of the conditions used in our study. However, even without these data, general conclusions about models of this type can still be made by examining several possible scenarios.

If the components are perceived unbiased, then the IOC rule prediction is that there will be *no bias* in the perceived direction of plaid motion. To obtain a plaid bias, there must be a component bias in at least one of the three required measurements. Furthermore, if the perceived direction and speed of each component grating are "co-biased" such that the resulting component velocities remain consistent with the unbiased constraint lines, and if the perceived orientation of the components is unbiased then the IOC rule predicts *no bias*. In this case, the computation recovers the original constraint lines and from these computes an unbiased plaid velocity. Other types of component biases will however produce biases in plaid direction.

To determine the potential effects of component biases on perceived plaid motion, it is necessary to know the biases in component speed, direction, and orientation. Unfortunately only one of these, the direction bias for gratings in elongated windows, has been examined, and it

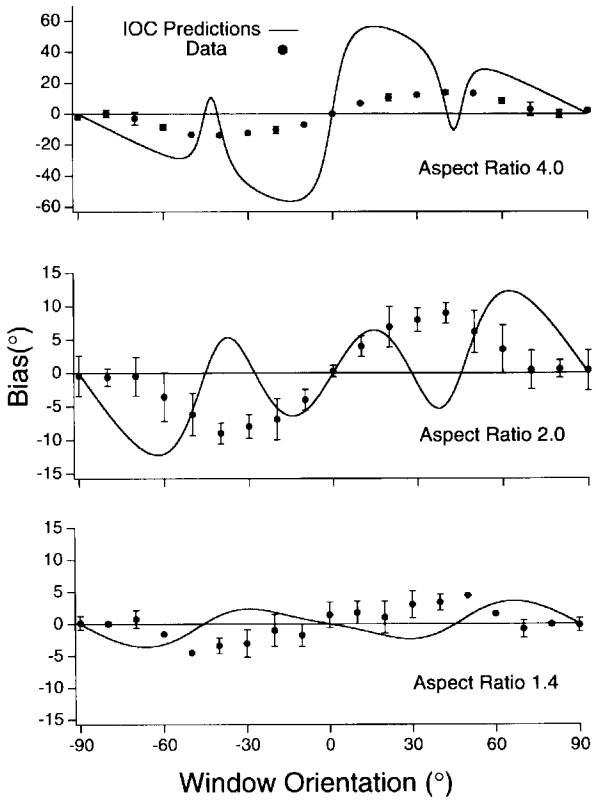

FIGURE 10. The biases in plaid direction predicted by the IOC model, which ignores the component orientation and uses only the biased component velocities, are plotted as a function of window orientation. Predictions are shown as the solid lines for the three aspect ratios: 1.4, 2.0, 4.0 (note the larger vertical range in the panel for aspect ratio of 4.0). For comparison, the filled circles show the average psychophysical data. The error bars represent the standard deviation over the two observers.

has been measured only in a limited number of conditions (Mulligan, 1991; Mulligan & Beutter, 1994). The dependence of the grating biases on the window aspect ratio is qualitatively similar to the predictions of a cross-correlation model, except that the measured biases decrease as spatial frequency increases while those predicted by cross-correlation are independent of spatial frequency. Therefore, we have used a cross-correlation model to estimate the component biases for our three aspect ratios, but have not calculated the predictions for varying the component spatial frequency.

We first computed the biased directions and speeds of gratings moving within elongated windows by maximizing the stimulus cross-correlation [Eq. (5), with the plaid replaced by a grating]. The results are shown in Fig. 9 for several aspect ratios and a grating spatial frequency of 0.6 c/deg. Although the model predicts large biases in both the component grating direction [Fig. 9(a)] and speed [Fig. 9(b)], these component biases are always linked: the biased speed and direction remain approximately consistent with the original grating constraint line [Fig. 9(b) solid lines]. Therefore, as pointed out above, if one assumes that the perceived orientations of the gratings are unaffected by the window, the biases in plaid direction predicted from these biased grating velocities are small [Fig. 7 and Fig. 8(a) dotted lines].

Another possibility is that the motion-processing system uses only the component directions and speeds as inputs to an IOC computation that implicitly assumes that the component orientation is orthogonal to the perceived component direction. This is equivalent to the orientation being misperceived as being in the direction perpendicular to the misperceived component direction of motion. The predicted plaid biases for this type of rule using the biased grating velocities from the cross-correlation computation (Fig. 9) are shown as solid lines in Fig. 10 for the three aspect ratios along with the average psychophysical results (filled circles). While this scenario does predict large biases in plaid direction, the patterns of biases do not even agree qualitatively with the data.

Each of the above IOC predictions depends on the method used to determine the component speed and direction biases. However, by focusing on the biases for

window orientations of 45 deg, it is possible to examine IOC models independently of any of our previous assumptions about component biases. Because our plaid components were always oriented ± 45 deg from the plaid direction, when the window orientation was 45 deg, one component was aligned with the major axis of the window, while the other component was aligned with the minor axis. When the components are aligned with the window, they are perceived to move in the direction perpendicular to their orientation, such that both their orientations and directions appear unbiased.* Therefore any bias in the perceived direction of plaid motion could only be caused by misperceptions of the component speeds. Our psychophysical data show that the average perceived biases for a window orientation of 45 deg and aspect ratios of 4.0, 2.0, and 1.4 are about 14.0, 8.2, and 4.4 deg, respectively. Using the IOC rule for perpendicular plaids, the relative speeds of the components necessary to produce these biases are $\tan(45\text{-}\theta_{bias})$, or 0.60, 0.75, and 0.86, respectively. Although it is possible that the elongated window may cause small biases in perceived component speed, it is doubtful that this effect could produce the required 14–40% change in relative speed. Because perceived grating speed depends on contrast (Thompson, 1982; Stone & Thompson, 1992; Muller & Greenlee, 1994), it is possible that component speed biases could result from small inequalities in the effective contrast of the components, that might be produced by the elongated window. However, any contrast effect produced by the window would be too small to be able to produce the large speed differences required to explain our data: the contrast ratio necessary to produce a ~25% change in perceived relative speed is ~7 (Stone & Thompson, 1992). Biases in component speed might also be caused by differences in the distances traveled (Brown, 1931). But because shorter paths produce faster perceived speeds, this cannot explain our results, since it would produce plaid biases in the wrong direction (towards the short axis).

In summary, we have examined several implementations of the IOC rule, and shown that none of them can predict our data:

1. If there is no bias in the perceived motion of the component grating, there is *no bias* in the predicted plaid direction.
2. If the component orientations are perceived veridically, and the perceived speed and direction of each component grating are co-biased in manner consistent with its constraint line, there is *no bias* in the predicted plaid direction.
3. If the component orientations are perceived veridically, and the perceived speed and direction of each component grating are those predicted by a cross-correlation model, the biases are *smaller* than our data.
4. If the perceived component orientations are ignored (or equivalently presumed orthogonal to the perceived direction of motion), and the constraints are computed by only using the component speed and direction predicted by a cross-correlation model, then the biases are *qualitatively different* than our data.
5. By examining the 45 deg window orientation case, we showed that our measured biases cannot be predicted by any IOC computation using plausible biases in the perceived component motions.

Therefore, the IOC model cannot explain our psychophysical data.

Vector-sum model

Wilson *et al.* (1992) have proposed a vector-sum motion processing model (hereafter referred to as the vector-sum model), which successfully predicts human perception of moving plaids under a variety of conditions (Ferrera & Wilson, 1990; Wilson *et al.*, 1992; Kim & Wilson, 1993; Wilson & Kim, 1994). Briefly, the vector-sum model has a Fourier pathway and a non-Fourier pathway. The Fourier pathway has an initial filtering stage whose output is passed through a Reichardt-like motion-energy computation which is followed by a gain-control mechanism. The non-Fourier pathway has an initial filtering stage whose output is squared, then processed by perpendicular filters which are tuned to lower spatial frequencies, and then is passed through a motion-energy mechanism followed by a gain-control mechanism. The output stage basically computes the average direction of the outputs of these two pathways weighted by the strengths of their responses.

We examined the predictions of this model by using a software implementation provided by Wilson. This implementation is designed for unwindowed plaid stimuli of infinite spatial extent. To apply the model to our finite stimuli, we calculated the response of its front-end filter† to the windowed component gratings of our stimuli and to identical unwindowed gratings. We computed the responses by centering the receptive field on the stimulus window and then integrating the stimulus weighted by the receptive-field filter over all space.‡ Each compo-

*Results from Mulligan (1991) and Mulligan and Beutter (1994) show that in the case where the grating is aligned with the long axis of the window, the perceived direction of motion is unbiased. Informal observations show that the perceived direction is also unbiased in the case where the grating is aligned with the short axis, and that, in both of these cases, the orientation appears unbiased.

†The spatial receptive fields in the vector sum model are defined as

$$F(x,y) = A\left[\exp\left(\frac{-x^2}{\sigma_1^2}\right) - B\cdot\exp\left(\frac{-x^2}{\sigma_2^2}\right)\right]\cdot\exp\left(\frac{-y^2}{\sigma_y^2}\right)$$

The parameter values are $\sigma_1 = 0.098$ deg, $\sigma_2 = 0.294$ deg, $A = 1022.7$, $B = 0.333$. Instead of specifying σ_y directly, the angular half-bandwidth at half height for the optimal spatial frequency, 1.7 c/deg was reported to be 22.5 deg. Using this, we calculated σ_y to be 0.405 deg.

‡Both the filter responses and the resultant predicted directions of motion depend on the position of the mechanisms relative to the center of the stimulus window. Avoiding the problem of determining how these different directions and speeds are integrated into a single percept, we have only calculated the predictions of a mechanism centered on the stimulus.

nent's effective contrast was then calculated as its actual contrast weighted by the ratio of the windowed-grating response to the unwindowed-grating response. We ran the simulations using these effective component contrasts as input, in place of the original contrast values.* The predicted biases for the 0.6 c/deg spatial frequency components and an aspect ratio of 2.0 are plotted as a function of the window angle in Fig. 7 as the dot-dashes. There are two major differences between the model predictions and the data. The predictions are both too small and of the wrong sign; instead of being biased toward the long axis of the window, the predicted biases are away from it. A similar pattern of results is predicted for other aspect ratios. The peak predicted biases are plotted as a function of aspect ratio in Fig. 8(a). The negative values indicate that the biases are in the wrong direction.

The vector-sum model predicts biases opposite those of our data because of the shape of its input spatial filters. To understand why this is so, it is again instructive to examine the case in which the window is oriented at 45 deg and the components move in the directions of the long and short axes of the window. The Fourier pathway responses contain two peaks of unequal magnitude because of the unequal effective component contrasts. The perhaps counter-intuitive result is that the responses to the grating moving in the direction of the short axis of the window are larger than those to the grating moving in the direction of the long axis of the window. This is caused by the shape of the input spatial filter: its spatial extent is larger in the direction perpendicular to the preferred direction of motion (a Gaussian with standard deviation 0.405 deg) than it is in the parallel direction (a difference of Gaussians whose standard deviations are 0.098 deg and 0.294 deg). Thus, since the predicted direction of motion is the weighted average of the grating responses, it is biased in the wrong direction, toward the short axis of the window. The effect of the non-Fourier pathway is small: it reduces the amount of the bias slightly. Because our stimuli are of equal spatial and temporal frequencies, the result of the non-Fourier squaring stage is a grating moving in the direction of plaid motion. (The squaring process produces gratings with frequencies that are the sum and difference of the component gratings' spatio-temporal frequencies. The sum moves in the plaid direction and the difference is stationary.) This produces non-Fourier pathway responses which are centered around the true plaid direction of motion. When included in the averaging they merely reduce the magnitude of the biases, but the biases are still always in the wrong direction.

*Our component gratings were 0.6 c/deg. Although the optimal frequency of the front end filters is 1.7 c/deg, documentation provided by Wilson states that the model operates well for spatial frequencies between 0.5 and 1.7 c/deg. To test whether the predictions depended critically on the plaid spatial frequency, we resimulated the model with our stimulus spatially rescaled (both the window and the plaid) so that its spatial frequency was 1.7 c/deg. The biases and trends we found for this 1.7 c/deg stimulus were similar to the results for our actual 0.6 c/deg stimulus.

We emphasize that the model's filter parameters are not arbitrary: they were empirically measured in a series of experiments (Wilson & Bergen, 1979; Wilson *et al.*, 1983; Phillips & Wilson, 1984; Wilson & Gelb, 1984) and used to model the psychophysical results of many plaid experiments (Ferrera & Wilson, 1990; Wilson *et al.*, 1992; Kim & Wilson, 1993; Wilson & Kim, 1994). Nonetheless we tuned the model parameters in an attempt to improve the match between the predictions and our data. Because, as discussed above, the filter shape causes the predicted biases to be in the direction opposite to those of our data, we examined the effects of varying the filter's height (σ_y). Experiments using drifting gratings (Anderson & Burr, 1991; Anderson *et al.*, 1991; Watson & Turano, 1995) have shown that the psychophysically measured receptive fields for moving stimuli have approximately equal height and width. We therefore reduced the input filter's height so that it was approximately equal to its width, and found that the predicted biases were still in the wrong direction, but they were smaller. This modification of the filter also had the undesirable effect of increasing the filter's angular bandwidth from 45 to 90 deg. To determine if an extreme modification of the filter might improve the predictions, we reduced the filter's height by an additional factor of 10. This manipulation did produce biases in the correct direction, but the biases were always < 1 deg, an order of magnitude smaller than our data. We conclude that even with drastic changes in its input filter shape, the vector-sum model cannot be modified to predict our data.

DISCUSSION

Our results show that the perceived direction of a plaid moving within an elongated window is biased toward the long axis of the window. The bias increases as the relative angle between the plaid direction and the long axis of the window increases, peaks at a relative angle of ~45 deg, and then decreases toward zero at 90 deg (when the plaid direction is aligned with the short axis of the window). This pattern of results was observed for all window aspect ratios and component spatial frequencies tested. The magnitude of the bias increases as the window aspect ratio increases and decreases as the plaid spatial frequency increases (or equivalently as the number of visible grating cycles increases). Although similar trends are present in the predictions of the motion-energy model, the cross-correlation model, and one of the modified IOC models, the magnitudes of the biases predicted by each of these models are much too small to account for our data. The predicted biases of the vector-sum model are in the wrong direction. Thus, none of these models can predict our psychophysical results.

Each of the above models responds to the "first-order" motion of the stimulus. Although our plaids are "first-order" stimuli and provide strong signals to each of the above models, it is possible that higher-order mechanisms, such as feature tracking, might also contribute to the percept [e.g. Lu & Sperling (1995)]. Possible features of our plaids include the bright or dark blobs and the

interblob regions. Because the position of these features is only slightly affected by the window shape, any feature tracking system will likely predict negligible biases. At present, no precise feature-tracking model exists for us to test explicitly, so quantitative estimation of the biases must await a specific proposal for how feature-tracking might be implemented. Nonetheless, it is doubtful that any such feature tracking system would generate biases large enough to explain our data. Furthermore, all of the models that we have examined predict biases which are too small. Therefore, we can also rule out strategies which use combinations of these processes, because they also would predict biases which are smaller than our data.

A number of previous studies have shown that varying factors such as contrast, spatial frequency, and adaptation state can produce significant biases in the perceived direction of plaid motion. If the component gratings have different contrasts, the direction of plaid motion is biased toward the direction of motion of the higher contrast grating (Stone *et al.*, 1990; Kooi *et al.*, 1992). If the component gratings have different spatial frequencies, the direction of plaid motion is biased toward the direction of motion of the grating of lower spatial frequency (Smith & Edgar, 1991; Kooi *et al.*, 1992). Derrington and Suero (1991) reported that adaptation to the direction of one of the components biased the perceived direction of plaid motion toward the direction of the other component. In these cases, the biases cannot be predicted by cross-correlation models, which are largely insensitive to these manipulations. However, these types of biases are not inconsistent with the predictions of modified motion-energy models or the vector-sum model. These data might also be predicted by an IOC model operating on misperceived component speeds, if the input speeds are modified to incorporate the known effects of contrast [e.g. Thompson (1982)], spatial frequency [e.g. Smith & Edgar (1990)], and adaptation [e.g. Thompson (1981); Muller & Greenlee (1994)]. Thus, these studies do not clearly distinguish between many of the leading motion processing models.

A different and perhaps more fundamental type of plaid direction misperception was reported by Ferrera and Wilson (1990). They examined the perceived direction of motion of Type II plaids, in which the plaid direction of motion lies outside the component directions. They found direction biases toward the components. This is not predicted by simple motion-energy or IOC models, but is predicted by the vector-sum model. However, interpreting these results is problematic. The appearance of these plaids is different than that of symmetric type I plaids, and at times Type II plaids appear not to move coherently. Because of this, Ferrera & Wilson (1987) originally called Type II plaids, "blobs", and noted that "(t)he motion of blobs does not always appear to be absolutely rigid, which might be taken to imply that coherence is not an all-or-none phenomenon, but that there may be cases of partial coherence" (p. 1788). This lack of coherence presents a problem for observers when asked to make a single direction judgment. On some trials, if observers perceive a partially coherent or even possibly incoherent stimulus and they are forced to make a direction judgment, they may respond to the motion of the components. If these responses are intermingled with those to trials in which the plaid is perceptually more coherent, a potential problem arises in interpreting the resultant psychometric data. The responses to the partially coherent trials will produce a bias toward the components, and also cause an increase in response variability which will result in a higher threshold [see footnote p. 1061 of Stone *et al.* (1990)]. This worrisome possibility cannot be distinguished from an actual increase in threshold and bias in perceived direction of motion. Thus it is impossible to rule out this explanation of Ferrera and Wilson's Type II plaid experiments, in which they find precisely this, both an increase in threshold (~6.5 deg for Type II compared to ~1 deg for symmetric Type I) and a bias toward the components' directions (~7.5 deg for Type II compared to ~0 deg for symmetric Type I).

To avoid the problem of coherence, we chose to use Type I symmetric, 90 deg, equal spatial and temporal frequency plaids, because they are known to cohere [Adelson & Movshon (1982); Welch & Bowne (1990); Smith & Edgar (1991); Smith (1992); Victor & Conte (1992); Kim & Wilson (1993); see however, Farid & Simoncelli (1994)]. Because of its simplicity, this type of plaid provides an extremely direct test of the basic principles of models. The components are identical except for their orientations, and thus issues of the interaction of different spatial and temporal frequencies are avoided. Additionally, the use of orthogonal components minimizes cross-orientation interactions. When the window is circularly symmetric, predicting its perceived direction of motion is particularly easy, almost any model (IOC, vector-sum, motion-energy, cross-correlation, even average direction) gets it right, but if the window is elongated, all of the models fail. Our data therefore provide a strong challenge to models of human motion perception, and suggest that there may be a basic problem in the way in which all of these models calculate the perceived velocity of moving patterns.

How then might the human brain estimate pattern velocity? The physiology and anatomy suggest a more elaborate approach in which spatial integration plays a key role. The stimulus is first processed by neurons with relatively local receptive fields which are orientation and spatial-frequency tuned (De Valois *et al.*, 1982a,b). At the next stage, motion is analyzed over more extended regions (Maunsell & Van Essen, 1983b; Albright, 1984; Mikami *et al.*, 1986b) and across a range of spatial frequencies (Movshon *et al.*, 1988) by neurons with larger receptive fields. Thus, it appears that motion information is first analyzed locally. These results are then combined across space and spatial scale to arrive at a unified global percept. In support of this type of processing, Kooi (1993) has shown that small local changes in a grating barberpole stimulus can change the way in which the aperture problem is resolved. He

showed that adding small indentations in the border of a rectangular aperture produces large changes in the global percept. To determine if spatial integration of local motion signals might also predict our plaid biases, we modified the cross-correlation and IOC models to only look at local stimulus patches. Not unexpectedly, we found that both the predicted directions and speeds varied significantly across space. Generally smaller direction and speed biases were predicted near the stimulus center and larger biases near the edges. To produce a unified percept of a drifting plaid, these local variations must be combined. Clearly the resultant pattern direction will depend on how this spatial integration is achieved. None of the models we examined address this issue. As in previous studies [e.g. Stone *et al.* (1990); Wilson *et al.* (1992)], our simulations were either global, a single measure across the whole visual field, or local, a single measure of the output of a sensor centered on the stimulus. However, expanding existing models to include an explicit spatial-integration rule may allow them to explain our results. In other words, it is not that we have shown that the algorithms tested above do not play a role in human motion processing but rather that, if they do, they will need to incorporate spatial integration before they can be used to predict human performance. Thus, our results show that even for a simple plaid stimulus, integration of motion signals across space may play a critical role in determining the perception of motion.

REFERENCES

Adelson, E. H. & Bergen, J. R. (1985). Spatiotemporal energy models for the perception of motion. *Journal of the Optical Society of America A, 2*, 284–299.

Adelson, E. H. & Movshon, J. A. (1982). Phenomenal coherence of moving visual patterns. *Nature, 300*, 523–525.

Albright, T. D. (1984). Direction and orientation selectivity of neurons in visual area MT of the macaque. *Journal of Neurophysiology, 52*, 1106–1130.

Albright, T. D., Desimone, R. & Gross, C. G. (1984). Columnar organization of directionally selective cells in visual area MT of the macaque. *Journal of Neurophysiology, 51*, 16–31.

Anderson, S. J. & Burr, D. R. (1991). Spatial summation properties of directionally selective mechanisms in human vision. *Journal of the Optical Society of America A, 8*, 1330–1339.

Anderson, S. J., Burr, D. R. & Morrone, M. C. (1991). Two-dimensional spatial and spatial-frequency selectivity of motion-sensitive mechanisms in human vision. *Journal of the Optical Society of America A, 8*, 1340–1351.

Beutter, B. R., Mulligan, J. B. & Stone, L. S. (1994a). The barberplaid illusion. *Investigative Ophthalmology and Visual Science, 35*(Suppl.), 2157.

Beutter, B. R., Mulligan, J. B. & Stone, L. S. (1994b). Direction of moving plaids is biased by asymmetric viewing windows. *Society for Neuroscience Abstracts, 20*, 772.

Britten, K. H., Shadlen, M. N., Newsome, W. & Movshon, J. A. (1993). Response of neurons in macaque MT to stochastic motion signals. *Visual Neuroscience, 10*, 1157–1169.

Brown, J. F. (1931). The visual perception of velocity. *Psychologische Forshungen, 14*, 199–232.

Bulthoff, H., Little, J. & Poggio, T. (1989). A parallel algorithm for real-time computation of optical flow. *Nature, 337*, 549–553.

Chubb, C. & Sperling, G. (1988). Drift-balanced random stimuli: A general basis for studying non-Fourier motion perception. *Journal of the Optical Society of America A, 5*, 1986–2007.

Chubb, C. & Sperling, G. (1989). Two-motion perception mechanisms revealed through distance-driven reversal of apparent motion. *Proceedings of the National Academy of Sciences, 86*, 2985–2989.

Derrington, A. M. & Badcock, D. R. (1992). Two-stage analysis of the motion of 2-dimensional patterns, what is the first stage? *Vision Research, 32*, 691–698.

Derrington, A. M. & Suero, M. (1991). Motion of complex patterns is computed from the perceived motions of their components. *Vision Research, 31*, 139–149.

De Valois, R. L., Albrecht, D. G. & Thorell, L. S. (1982a). Spatial frequency selectivity of cells in macaque visual cortex. *Vision Research, 22*, 545–559.

De Valois, R. L., Yund, E. W. & Hepler, N. (1982b). The orientation and direction sensitivity of cells in macaque visual cortex. *Vision Research, 22*, 531–544.

De Yoe, E. A. & Van Essen, D. C. (1985). Segregation of efferent connections and receptive field properties in visual area V2 of the macaque. *Nature, 32*, 58–61.

Farid, H. & Simoncelli, E. P. (1994). The perception of transparency in moving square-wave plaids. *Investigative Ophthalmology and Visual Science, 35*(Suppl.), 1271.

Fennema, C. L. & Thompson, W. B. (1979). Velocity determination in scenes containing several moving objects. *Computer Graphics and Image Processing, 9*, 301–315.

Ferrera, V. P. & Wilson, H. R. (1987). Direction specific masking and the analysis of motion in two dimensions. *Vision Research, 27*, 1783–1796.

Ferrera, V. P. & Wilson, H. R. (1990). Perceived direction of moving two-dimensional patterns. *Vision Research, 30*, 273–287.

Heeger, D. J. (1987). Model for the extraction of image flow. *Journal of the Optical Society of America A, 4*, 1455–1471.

Hubel, D. H. & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology, 195*, 215–243.

Kim, J. & Wilson, H. R. (1993). Dependence of plaid motion coherence on component grating directions. *Vision Research, 33*, 2479–2489.

Kooi, F. L. (1993). Local direction of edge motion causes and abolishes the barberpole illusion. *Vision Research, 33*, 1347–1351.

Kooi, F. L., De Valois, K. K., Grosof, D. A. & De Valois, R. L. (1992). Properties of the recombination of one-dimensional motion signals into a pattern motion signal. *Perception and Psychophysics, 52*, 415–424.

Leese, J. A., Novak, C. S. & Taylor, V. R. (1970). The determination of cloud pattern motion from geosynchronous satellite image data. *Pattern Recognition, 2*, 279–292.

Lu, Z. L. & Sperling, G. (1995). The functional architecture of human visual motion perception. *Vision Research, 35*, 2697–2722.

Maunsell, J. H. R. & Van Essen, D. C. (1983a). The connections of the middle temporal area (MT) and their relationship to cortical hierarchy in the macaque monkey. *Journal of Neuroscience, 3*, 2563–2586.

Maunsell, J. H. R. & Van Essen, D. C. (1983b). Functional properties of neurons in the middle temporal visual of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation. *Journal of Neurophysiology, 49*, 1127–1147.

Mikami, A., Newsome, W. T. & Wurtz, R. H. (1986a). Motion selectivity in macaque visual cortex. I. Mechanisms of direction and speed selectivity in extrastriate area MT. *Journal of Neurophysiology, 55*, 1308–1327.

Mikami, A., Newsome, W. T. & Wurtz, R. H. (1986b). Motion selectivity in macaque visual cortex. II. Spatiotemporal range of directional interactions in MT and V1. *Journal of Neurophysiology, 55*, 1328–1339.

Movshon, J. A., Adelson, E. H., Gizzi, M. S. & Newsome, W. T. (1986). The analysis of moving visual patterns. *Experimental Brain Research, 11*(Suppl.), 117–151.

Movshon, J. A. & Newsome, W. T. (1984). Functional characteristics of striate cortical neurons projecting to MT in the macaque. *Society for Neuroscience Abstracts, 10*, 983.

Movshon, J. A., Newsome, W. T., Gizzi, M. S. & Levitt, J. B. (1988). Spatio-temporal tuning and speed sensitivity in macaque visual

cortical neurons. *Investigative Ophthalmology and Visual Science, 29*(Suppl.), 327.

Muller, R. & Greenlee, M. W. (1994). Effect of contrast and adaptation on the perception of the direction and speed of drifting gratings. *Vision Research, 34*, 2071–2092.

Mulligan, J. B. (1986). Minimizing quantization errors in digitally controlled CRT displays. *Color Research and Application, 11*, S47–S51.

Mulligan, J. B. (1991). A continuous version of the barber-pole illusion. *Investigative Ophthalmology and Visual Science, 32*(Suppl.), 829.

Mulligan, J. B. & Beutter, B. R. (1994). Additional twists on a continuous barber-pole illusion. *Investigative Opthamology and Visual Science, 35*(Suppl.), 2157.

Mulligan, J. B. & Stone, L. S. (1989). Halftoning method for the generation of motion stimuli. *Journal of the Optical Society of America A, 6*, 1217–1227.

Newsome, W. T., Mikami, A. & Wurtz, R. H. (1986). Motion selectivity in macaque visual cortex. III. Psychophysics and physiology of apparent motion. *Journal of Neurophysiology, 55*, 1340–1351.

Phillips, G. C. & Wilson, H. R. (1984). Orientation bandwidths of spatial mechanisms measured by masking. *Journal of the Optical Society of America A, 1*, 226–232.

Power, R. P. & Mouldon, B. (1992). Spatial gating effects on judged motion of gratings in apertures. *Perception, 21*, 449–463.

Rodman, H. R. & Albright, T. D. (1989). Single-unit analysis of pattern-motion selective properties in the middle temporal visual area (MT). *Experimental Brain Research, 75*, 53–64.

Rodman, H. R., Gross, C. G. & Albright, T. D. (1989). Afferent basis of visual response properties in area MT of the macaque. I. Effects of striate cortex removal. *Journal of Neuroscience, 9*, 2033–2050.

Rodman, H. R., Gross, C. G. & Albright, T. D. (1990). Afferent basis of visual response properties in area MT of the macaque. II. Effects of superior colliculus removal. *Journal of Neuroscience, 10*, 1154–1164.

van Santen, J. P. H. & Sperling, G. (1984). Temporal covariance model of human motion perception. *Journal of the Optical Society of America A, 1*, 451–473.

van Santen, J. P. H. & Sperling, G. (1985). Elaborated Reichardt detectors. *Journal of the Optical Society of America A, 2*, 300–320.

Smith, A. T. (1992). Coherence of plaids comprising components of disparate spatial frequencies. *Vision Research, 32*, 393–397.

Smith, A. T. & Edgar, G. K. (1990). The influence of spatial frequency on perceived temporal frequency and perceived speed. *Vision Research, 30*, 1467–1474.

Smith, A. T. & Edgar, G. K. (1991). Perceived speed and direction of complex gratings and plaids. *Journal of the Optical Society of America A, 8*, 1161–1171.

Stone, L. S. (1990). Precision in the perception of plaid motion. *Investigative Ophthalmology and Visual Science, 31*(Suppl.), 172.

Stone, L. S. & Thompson, P. (1992). Human speed perception is contrast dependent. *Vision Research, 32*, 1535–1549.

Stone, L. S., Watson, A. B. & Mulligan, J. B. (1990). Effect of contrast on the perceived direction of a moving plaid. *Vision Research, 30*, 1049–1067.

Thompson, P. (1981). Velocity after-effects: The effects of adaptation to moving stimuli on the perception of subsequently seen stimuli. *Vision Research, 21*, 337–345.

Thompson, P. (1982). Perceived rate of movement depends on contrast. *Vision Research, 22*, 377–380.

Ungerleider, L. G. & Desimone, R. (1986). Cortical connections of area MT in the macaque. *Journal of Comparative Neurology, 248*, 190–222.

Ungerleider, L. G., Desimone, R., Galkin, T. W. & Mishkin, M. (1984). Subcortical projections of area MT in the macaque. *Journal of Comparative Neurology, 223*, 368–386.

Victor, J. D. & Conte, M. M. (1992). Coherence and transparency of moving plaids composed of Fourier and non-Fourier gratings. *Perception and Psychophysics, 52*, 403–414.

Wallach, H. (1935). Uber visuell wahrgenommene Bewegungsrichtung. *Psychologische Forschungen, 20*, 325–380.

Watson, A. B. & Ahumada, A. J. (1983). A look at motion in the frequency domain. In *Motion: Perception and representation*. New York: Association for Computing Machinery (also published as NASA TM-84355).

Watson, A. B. & Ahumada, A. J. (1985). Model of human visual-motion sensing. *Journal of the Optical Society of America A, 2*, 322–341.

Watson, A. B., Thompson, P. G., Murphy, B. J. & Nachmias, J. (1980). Summation and discrimination of gratings moving in opposite directions. *Vision Research, 20*, 341–347.

Watson, A. B. & Turano, K. (1995). The optimal motion stimulus. *Vision Research, 35*, 325–336.

Welch, L. (1989). The perception of moving plaids reveals two motion-processing stages. *Nature, 337*, 734–736.

Welch, L. & Bowne, S. F. (1990). Coherence determines speed discrimination. *Perception, 19*, 425–435.

Wilson, H. R. & Bergen, J. R. (1979). A four-mechanism model for threshold spatial vision. *Vision Research, 19*, 19–32.

Wilson, H. R., Ferrera, V. P. & Yo, C. (1992). A psychophysically motivated model for two-dimensional motion perception. *Visual Neuroscience, 9*, 79–97.

Wilson, H. R. & Gelb, D. J. (1984). Modified line-element theory for spatial-frequency and width discrimination. *Journal of the Optical Society of America A, 1*, 124–131.

Wilson, H. R. & Kim, J. (1994). Perceived motion in the vector sum direction. *Vision Research, 34*, 1835–1842.

Wilson, H. R., McFarLane, D. K. & Phillips, G. C. (1983). Spatial-frequency tuning of orientation selective units estimated by oblique masking. *Vision Research, 23*, 873–882.

Zeki, S. M. (1974). Functional organization of a visual area in the posterior bank of the superior temporal sulcus of the rhesus monkey. *Journal of Physiology (London), 236*, 546–573.

Acknowledgements—The authors would like to thank Hugh Wilson for kindly providing a MATLAB® implementation of the vector-sum motion model (Wilson *et al.*, 1992). We would also like to thank Andrew Watson and Preeti Verghese for their helpful comments on earlier drafts and the NASA Ames Vision Group for their general support. This work was supported by a NASA RTOP 199-16-12-37 to LS, and by a National Research Council postdoctoral associateship to BB.